arXiv:1407.5374v7 [cs.DM] 13 May 2015

On the Algorithmic Lovász Local Lemma and Acyclic Edge Coloring*†

Ioannis Giotis
Department of Mathematics, National & Kapodistrian University of Athens
and Computer Technology Institute & Press "Diophantus"
igiotis@cs.upc.edu

Lefteris Kirousis
Department of Mathematics, National & Kapodistrian University of Athens
and Computer Technology Institute & Press "Diophantus"
lkirousis@math.uoa.gr

Kostas I. Psaromiligkos
Department of Mathematics, National & Kapodistrian University of Athens
kostaspsa@gmail.com

Dimitrios M. Thilikos
Department of Mathematics, National & Kapodistrian University of Athens,
Computer Technology Institute & Press "Diophantus"
and AlGCo project-team, CNRS, LIRMM, Montpellier
sedthilk@thilikos.info

May 14, 2015

Abstract

The algorithm for Lovász Local Lemma by Moser and Tardos gives a constructive way to prove the existence of combinatorial objects that satisfy a system of constraints. We present an alternative probabilistic analysis of the algorithm that does not involve reconstructing the history of the algorithm. We apply our technique to improve the best known upper bound to acyclic chromatic index. Specifically we show that a graph with maximum degree Δ has an acyclic proper edge coloring with at most $\lceil 3.74(\Delta-1)\rceil+1$ colors, whereas the previously known best bound was $4(\Delta-1)$. The same technique is also applied to improve corresponding bounds for graphs with bounded girth. An interesting aspect of the latter application is that the probability of the "undesirable" events do not have a uniform upper bound, i.e., it constitutes a case of the *asymmetric* Lovász Local Lemma.

*A preliminary version of this paper was presented at the SIAM Meeting on Analytic Algorithmics and Combinatorics (ANALCO), 2015.

†Research co-financed by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF), Research Funding Program: ARISTEIA II.

1 Introduction

The Lovász Local Lemma (LLL) first appeared in 1975 in a paper by Erdős and Lovász [6].

> "The proof is so elementary that it could, and I think it should, be taught in a first course in probability. It has had and continues to have a profound effect on probabilistic method."

J. Spencer wrote about it in the book on his Durango lectures [20]. This elegant proof however was non-constructive. Up until the seminal work by Moser [16] and Moser and Tardos [17], and despite important work on the question of supplying a constructive proof by Beck [4], Alon [2], Srinivasan [22] and others, the question and its answer succinctly given in the book on the Durango lectures [20]: "Algorithm? Sometimes!" remained valid.

The constructive approaches of Moser [16] and Moser and Tardos [17] showed how to algorithmically produce a satisfying assignment to a Constraint Satisfaction Problem (CSP), given that the variables are independent. The randomized algorithms in both papers look for an unsatisfied constraint and resample its variables, until no unsatisfied constraint is left. In Moser [16], the search is local, in the sense that the algorithm looks for unsatisfied constraints in the neighborhood of the current one, whereas in [17] there is no locality restriction for the search[1]. In both papers the expected time of steps for such an algorithm to halt is shown to be linear in the number of variables, given that the degree of the dependency graph of the constraints does not exceed a certain constant fraction of the inverse of the probability of a constraint not to be satisfied. The proofs depend on supplying a witness structure which encodes the history of the random choices of the algorithm. For an elegant presentation of the Moser and Tardos' proof [17] see Spencer [21]. Another proof approach (closer to the spirit of the proof in [16]) depends on the fact that the size of the witness structure cannot be less than the entropy of the sequence of random choices made by the algorithm. See [24] for a short exposition of the basic idea of this information theoretic argument known as "*entropy compression argument*". For a recent survey, see Szegedy [23].

In this work, we work with Moser's original algorithm in [16] and we use the Principle of Deferred Decisions of Knuth [14] to show that resampling produces an assignment with the same distribution as the original one (conditional that the event that initiated this particular resampling is satisfied). Thus we manage to express the probability that at least n steps are needed until the algorithm halts by a simple recurrence relation. This recurrence relation is then asymptotically analyzed by classical tools of the Analysis of Algorithms. It turns out that given that the degree of the dependency graph does not exceed a certain constant fraction of the inverse probability of an undesirable event to occur, the probability that the algorithm lasts for n steps is exponentially small in n, after a cutoff point (a result made known to us by Achlioptas and Iliopoulos [1], but through a proof based on the entropic method). We believe that this direct and completely elementary probabilistic proof avoids some of the intricacies of

[1] The algorithm in [16] was presented and analyzed only for the SAT problem whereas the one in [17] was presented and analyzed in a general setting; the generalization of the algorithm in [16] for arbitrary events is straightforward, however the analysis in [16] does not immediately generalize.

the entropy compression argument. It also completely unveils, we think, the extreme elegancy of Moser's algorithm.

We used our alternative approach to a particular coloring problem, namely the acyclic edge coloring. The entropy compression method had been used by Esperet and Parreau [7] to find bounds for the acyclic chromatic index, the least number of colors needed to properly color the edges of a graph so that no bichromatic cycle exists. Our approach leads to a direct and simple probabilistic analysis that yields the upper bound of $\lceil 3.74(\Delta-1)\rceil+1$ for the acyclic chromatic index, improving the bound of $4(\Delta-1)$ given by Esperet and Parreau. We also improve their bounds for graphs with bounded girth.

There are two interesting remarks to be made with respect to the application of our technique to the edge coloring problem:

- There is no uniform upper bound neither for the probabilities of the events considered as "undesirable" nor for the neighborhood size of the dependency graph. So the application in this case is essentially an instance of the Asymmetric Lovász Local Lemma.
- The variables that determine the "undesirable" events, namely the edges of the graph to be colored, are not independent.

A proof of the *asymmetric* Lovász Local Lemma by the same technique has been given recently by Giotis et al. [11]. Also recently, Harvey and Vondrák [13] gave a proof of the lopsided asymmetric Lovász Local Lemma by a similar technique (the events in that work are over an arbitrary probability space, as also in Kolipaka and Szegedy [15]).

2 Algorithmic Lovász Local Lemma

Let $X_i, i = 1,\ldots,l$ be mutually independent random variables on a common probability space, taking values in the sets $D_i, i = 1,\ldots,l$, respectively (we opted for the letter l to stand for the number of variables, as by n we will denote number of algorithmic steps).

Let $E_j, j = 1,\ldots,m$ be a sequence of events, each depending on a sequence of the random variables X_i. The sequence of variables that an event E_j depends on is called the *scope* of E_j and is denoted by e^j. Events are assumed to be ordered as in the sequence (i.e., according to their index). The events E_j are considered "undesirable", i.e., the objective is to design a randomized algorithm that will return an assignment α to the variables X_i for which none of the events E_j hold.

We say that two events overlap, and write $E_j \sim E_k$, if $e^j \cap e^k \neq \emptyset$. The binary, reflexive and symmetric binary relation $\sim$ defines a graph with all vertices looped (but no multiple edges) called the *dependency graph* of the events.

For $j = 1,\ldots,m$, let N_j be the neighborhood of the event E_j in the dependency graph, i.e., $N_j = \{E_k \mid E_k \sim E_j\}$ (observe that $E_j \in N_j$).

Let Δ be the maximum of the cardinalities $|N_j|$ (i.e., the max degree of the dependency graph counting the loop as contributing 1 to the degree) and let p be the max of the probabilities $\Pr[E_j]$.

M-ALGORITHM

1: Sample the variables $X_i, i = 1, \ldots, l$ and let α be the resulting assignment of values to them.
2: **while** there exist an event that occurs under the current assignment, let E_j be the least indexed such event **do**
3: RESAMPLE(E_j)
4: **end while**
5: Output current assignment α

RESAMPLE *(E_j)*

1: Resample the variables in the scope e^j.
2: **while** some $E_k \in N_j$ occurs under the current assignment, let E_k be the least indexed such event **do**
3: RESAMPLE(E_k)
4: **end while**

Figure 1: *Moser's algorithm*

Theorem 1 (Lovász Local Lemma)**.** *If $ep\Delta \leq 1$, then $Pr[\overline{E_1} \wedge \overline{E_2} \wedge \cdots \wedge \overline{E_m}] > 0$, i.e., there exists an assignment to the variables X_i for which none of the events E_i hold (by e we denote the base of the natural logarithm, i.e., $e = 2.71828...$).*

The original proof of Theorem 1, first presented essentially in this form in [19], was non-constructive, but was given for arbitrary events, i.e., without the assumption that the events depended on independent random variables. Below, we will give an algorithmic proof Theorem 1 within the framework already described, i.e., assuming the dependence of the events on independent variables.

We first present our algorithm in Figure 1, which is a direct derivation of the one given by Moser in [16].

It is trivial to see that if M-ALGORITHM ever stops, it returns an assignment for which none of the "undesirable" events $E_1, \ldots, E_m$ holds.

We first prove:

Lemma 1. *Consider an arbitrary call of* RESAMPLE *(E_j). Let $\mathcal{X}$ be the set of variables that do not belong to the scope of an event that occurs just when this call is to start. Then, if the call terminates, variables in $\mathcal{X}$ will also not belong to the scope of an event occurring when this call is just completed.*

Proof. Suppose some variable X_i in $\mathcal{X}$ appears in the scope of an event E_s that occurs during the call RESAMPLE (E_j). This implies that for some event E_k sharing a variable with E_s, the call RESAMPLE (E_k) was made during the execution of RESAMPLE (E_j) —either as a recursive call within RESAMPLE (E_j) or as RESAMPLE (E_j) itself— and E_s occurred after a resampling of the variables in e^k. Let $E_{k_{\text{last}}}$ be the last such event E_k, i.e., the last event during the execution of RESAMPLE (E_j) for which there is an event E_s such that $E_k \sim E_s$ and X_i is in e^s and E_s occurs after a resampling of the variables e^k. Let also $E_{s_{\text{last}}}$ be the last such E_s for $E_{k_{\text{last}}}$. Because of the "while" loop of line 2 of RESAMPLE, the call RESAMPLE *($E_{k_{\text{last}}}$)* will not terminate until all the events

in the neighborhood $N_{k_{\text{last}}}$ of $E_{k_{\text{last}}}$, and therefore $E_{s_{\text{last}}}$ itself, do not occur. Assuming RESAMPLE (E_j) terminates, RESAMPLE $(E_{k_{\text{last}}})$ must also terminate and therefore just at the end of RESAMPLE $(E_{k_{\text{last}}})$, $E_{s_{\text{last}}}$ does not occur. □

Let us now call *root call* of RESAMPLE a call of RESAMPLE made when executing line 3 of M-ALGORITHM. We refer to calls of RESAMPLE from within another call of RESAMPLE as *recursive* calls. On a particular execution of our algorithm, let us call *phase* the execution period from the start until the completion of a root call of RESAMPLE.

By Lemma 1, we know that the events of root calls of RESAMPLE have pairwise non-intersecting scopes and in particular they are pairwise distinct. Therefore:

Corollary 1. *There are at most m phases in any execution of* M-ALGORITHM.

Our aim in the sequel will be to bound from above, as a function of n, the probability that M-ALGORITHM takes at least n steps. As algorithmic steps we count only calls of RESAMPLE (root or recursive). Towards finding a bound for the probability that n steps are needed, we will argue that the probability of an event occurring at any step of the algorithm, conditional that the previous steps have already been performed to completion, is bounded by p, the bound of the probability that an event occurs when the algorithm starts.

In the sequel, by *rooted forest* we will mean a forest of trees such that each tree has a special node designated as its root.

The nodes of rooted forests are labeled by the events $E_j, j = 1, \ldots, m$ with repetitions of the labels allowed.

Definition 1. *A labeled rooted forest is called* feasible *if*

i. The labels of the roots of the trees have pairwise disjoint scopes,

ii. the labels of the children of every node have pairwise disjoint scopes and

iii. the label of each of the children of every node u shares a variable (is a neighbor in the sense of $\sim$) with the label of u.

Apart form the restrictions between labels imposed by the definition, the label of a node u can share a variable or even coincide with the label of another node v (even if u, v belong to different trees), as long as u and v are not siblings and are not both roots.

The nodes of a feasible forest are ordered as follows: children of the same node are ordered as their labels are; nodes in the same tree are ordered by preorder (respecting the ordering between siblings) and finally if the label on the root of a tree T_1 precedes the label of the root of T_2, all nodes of T_1 precede all nodes of T_2.

The number of nodes of a forest F is denoted by $|F|$.

Now consider the first n consecutive steps of a run of M-ALGORITHM. We construct in a unique way, depending on the steps, a feasible forest with n nodes as follows:

i. The forest under construction will have as many roots as root calls of RESAMPLE exist within the n steps. These roots will be labeled by the event of the corresponding root call.

ii. A tree that corresponds to a root call RESAMPLE (E_j) will have as many non-root nodes as the number of recursive calls of RESAMPLE within RESAMPLE (E_j) (that appear in the sequence of the given n consecutive steps). The non-root nodes will be labeled by the events of these recursive calls.

iii. The non-root nodes are organized within the tree with root labeled with E_j so that a node that corresponds to a call RESAMPLE (E_k) is parent to a root that corresponds to a call RESAMPLE (E_l) if RESAMPLE (E_l) ever appears immediately on top of RESAMPLE (E_k) in the recursive stack that implements the root call RESAMPLE (E_j).

It is straightforward to check, inspecting the succession of steps of M-ALGORITHM in any run, and making use of Lemma 1, that a forest constructed as above from a number of consecutive steps of a run of M-ALGORITHM is indeed a feasible forest in the sense of Definition 1. It is not true however that every feasible forest can be constructed as above from a number of consecutive of steps of some run of M-ALGORITHM. For example a feasible forest where two children of a node are labeled by events E_j and E_k, respectively, such that E_j occurs if and only if E_k does, can never be constructed as above.

Definition 2. *The* witness *forest of the first n steps of a run of* M-ALGORITHM *is the feasible forest constructed as described above. Given a feasible forest F with n nodes, we denote by W_F the event that* M-ALGORITHM *runs for at least n steps and during these n steps, F is constructed. We also set:*

$$\hat{P}_n = \Pr\left[\bigcup_{F:\,|F|=n} W_F\right].$$

N.B. The event W_F is none of the events E_j, it is rather an event that refers to which events E_j and in what order the RESAMPLE calls of the randomized M-ALGORITHM are applied.

Obviously the probability that M-ALGORITHM takes at least n steps is equal to $\hat{P}_n$, notationally:

$$\Pr[\text{ M-ALGORITHM takes at least } n \text{ steps}] = \hat{P}_n. \tag{1}$$

Towards finding an upper bound for $\hat{P}_n$, we now describe a randomized algorithm, to be called validation algorithm, notationally VALALG, that takes as input a feasible forest F with n nodes, labeled with the events $E_{j_1}, \ldots, E_{j_n}$ (ordered as their respective nodes) and outputs a Boolean value `success` or `failure`. The algorithm is given in Figure 2.

Intuitively, with input a feasible forest with labels $E_{j_1}, \ldots, E_{j_n}$, VALALG initially generates a random sampling of the variables, then checks the current event and, if it holds, resamples its variables and goes to the next event. If the algorithm manages to go through *all* events, it returns `success`, otherwise, at the first event that does not hold under the current assignment, it returns `failure` and stops.

We first give a result about the probability distribution of the assignment to the variables X_i during the execution of VALALG.

ValAlg (forest F with labels on the roots $E_{j_1} \ldots, E_{j_n}$)

1: Sample the variables $X_i, i = 1, \ldots, l$.
2: **for** $s = 1, \ldots, n$ **do**
3: **if** E_{j_s} does not occur under current assignment **then**
4: **return** `failure` and exit.
5: **else**
6: Resample the variables of E_{j_s}.
7: **end if**
8: **end for**
9: **return** `success`.

Figure 2: *The validation algorithm*

Lemma 2 (Randomness lemma)**.** *At every instant during the execution of* ValAlg *that immediately follows a successful (i.e., no* `failure` *returned) iteration of the* **for** *loop of Step 2, the distribution of the current assignment of values to the variables $X_i, i = 1, \ldots, l$ is as if* all *variables have been sampled anew. Therefore the probability of any event occurring at this instant is bounded by p.*

Proof. The result immediately follows from the Principle of Deferred Decisions of Knuth [14]. Indeed the variables of the current assignment that belong to the scope of the last checked event have been actually sampled anew, whereas for the others we may assume that are sampled anew by the Principle of Deferred Decisions. □

N.B. The above result does not hold for M-Algorithm, because that algorithm exposes information also for events other than the one whose variables are resampled (e.g. that the all events preceding the one whose variables are resampled do not occur). This is the reason ValAlg is introduced.

Definition 3. *Given a feasible forest F with n nodes, we say that F is validated by* ValAlg *if the latter returns* `success` *on input F. The event of this happening is denoted by V_F. We also set:*

$$P_n = \Pr\left[\bigcup_{F:\,|F|=n} V_F\right].$$

Observe that by union bound:

$$P_n \leq \sum_{F:\,|F|=n} \Pr[V_F]. \tag{2}$$

We now claim:

Lemma 3. *For any feasible forest F, the event W_F implies the event V_F, therefore:*

$$\hat{P}_n \leq P_n. \tag{3}$$

Proof. Let $E_{j_1}, \ldots, E_{j_n}$ be the labels of F (in the order of the corresponding nodes). Consider a run of n steps of M-ALGORITHM that gives F as its witness forest and let $\alpha_1 \ldots, \alpha_n$ be the sequence of consecutive assignments generated by this run. Then we get immediately by inspecting M-ALGORITHM that:

- α_1 is obtained by independently sampling the variables $X_i, i = 1, \ldots, l$.
- For $s = 1, \ldots n$, E_{j_s} does not hold under α_s.
- For $s = 1, \ldots n-1$, α_{s+1} is obtained from α_s by resampling the variables of E_{j_s}.

Therefore the sequence $\alpha_1 \ldots, \alpha_n$ comprises the assignments of a run of VALALG that on input F, returns `success`. □

It is the probability that a feasible forest F with n nodes is validated that we will bound from above in the sequel, thus by Lemma 3 achieving to bound from above the probability of getting a witness forest with n nodes and in turn, by equality (1), to bound from above the probability of M-ALGORITHM taking at least n steps.

Now let F be a feasible forest with trees $T_1, \ldots, T_m$ with roots labeled with $E_1, \ldots, E_m$, respectively.

Note that we have implicitly assumed that all events $E_j, j = 1, \ldots, m$ appear as labels of the root of some tree of F; if this is not the case, that is if some E_j does not appear as root of a tree of F, set $T_j = \emptyset$. Also we will assume that for all nodes u there are exactly Δ subtrees hanging from u by taking as subtree the empty tree for missing children.

Then each tree T_j can be considered as a labeled feasible forest with a single tree, so it makes sense to consider $\Pr[V_{T_j}], j = 1 \ldots, m$ (if $T_j = \emptyset$, we clearly assume that $\Pr[V_{T_j}] = 1$) An immediate corollary of Lemma 2 is the following:

Lemma 4.

$$\Pr[V_F] \leq \prod_{j=1}^{m} \Pr[V_{T_j}].$$

Also if a single-tree forest T is comprised of the trees $T_1, \ldots T_\Delta$ hanging from the root, then:

$$\Pr[V_T] \leq p \prod_{k=1}^{\Delta} \Pr[V_{T_k}].$$

Proof. Indeed for the first inequality observe that by Lemma 2, the assignment α that VALALG produces after it has examined the last event of T_{j-1}, conditional that `success` was returned, is distributed as if α was obtained by sampling anew all variables.

For the second part, first notice that the factor p appears because the event in the root of T should occur when T is validated. Then, in a similar way as the first part, notice that the assignments that exist at the start of $T_1, \ldots, T_\Delta$ during the execution of VALALG are all distributed as if all variables were sampled anew. □

Let now $\mathcal{F}_{E_j}$ be the collection of feasible forests each of which comprises of a *single* tree whose root is labeled with E_j. Assume that numbers Q_n are given such that such that for every n and every event $E_j, j = 1, \ldots, m$:

$$\sum_{\substack{T:\, T\in\mathcal{F}_{E_j} \\ \&|T|=n}} \Pr[V_T] \le Q_n. \tag{4}$$

By rearranging sums, from inequality (2) and the first inequality of Lemma 4, we get that:

$$P_n \le \sum_{\substack{n_1+\cdots+n_m=n \\ n_1,\ldots,n_m\ge 0}} Q_{n_1}\cdots Q_{n_m}. \tag{5}$$

Also, from the second inequality of Lemma 4 we conclude that numbers Q_n satisfying (4) can be defined by the recurrence:

$$Q_n = \begin{cases} p\left(\sum\limits_{\substack{n_1+\cdots+n_\Delta=n-1 \\ n_1,\ldots,n_\Delta\ge 0}} Q_{n_1}\cdots Q_{n_\Delta}\right) & \text{if } n\ge 1, \\ 1 & \text{if } n=0. \end{cases} \tag{6}$$

Lemma 5 (Phase bound)**.** *Q_n is asymptotically bounded from above by*

$$\sqrt{1+\frac{1}{\Delta-1}}\left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1}p\Delta\right)^n < \sqrt{1+\frac{1}{\Delta-1}}\,(ep\Delta)^n.$$

Proof. Let $Q(z)$ be the OGF of Q_n. Multiply both its sides of (6) with z^n and sum for $n\ge 1$ to get (since $Q_0=1$):

$$Q(z)-1 = zpQ(z)^\Delta \tag{7}$$

Let now $W=W(z)=Q(z)-1$. Then by (7), we have $W=zp(W+1)^\Delta$.

Let now $\phi(W)=p(W+1)^\Delta$. Then Equation (7) is equivalent to

$$W(z) = z\phi(W(z)). \tag{8}$$

Now apply Lagrange's Inversion Formula (see e.g. [9, Theorem A.2]) to get that for $n\ge 1$:

$$\begin{aligned} Q_n = [z^n]W &= (1/n)[u^{n-1}](\phi(u))^n = (1/n)p^n\binom{\Delta n}{n-1} \\ &= \frac{1}{(\Delta-1)n+1}p^n\binom{\Delta n}{n}. \end{aligned}$$

Therefore by Stirling approximation there is a constant $c'>0$ (depending on

Δ) such that:

$$\begin{aligned} Q_n <& e^{c'/n} \frac{1}{(\Delta-1)n+1} \frac{1}{\sqrt{2\pi n}} \\ & \sqrt{1+\frac{1}{\Delta-1}} \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^n \\ <& e^{c'/n} \sqrt{1+\frac{1}{\Delta-1}} \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^n \\ <& e^{c'/n} \sqrt{1+\frac{1}{\Delta-1}} \left(ep\Delta\right)^n . \end{aligned}$$

This concludes the proof of Lemma 5. □

Lemma 6 (Algorithm bound)**.** *There is a constant $A > 1$ (depending on Δ) such that P_n is bounded from above by*

$$(An)^m \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^n < (An)^m \left(ep\Delta\right)^n .$$

Proof. By (5) and Lemma 5 we have that for some $A > 1$ (depending on Δ), P_n is bounded from above by

$$\begin{aligned} \sum_{\substack{n_1+\cdots+n_m=n \\ n_1,\ldots,n_m\geq 0}} & A \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^{n_1} \cdots \\ & \cdots A \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^{n_m} \\ & \leq (An)^m \left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right)^n . \end{aligned}$$

This concludes the proof of Lemma 6. □

Now to have that P_n becomes exponentially small in n when

$$\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta < 1,$$

we must have that

$$m\log n + m\log A + n\log\left(\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta \right) < 0.$$

Therefore by equality (1) and inequality (3) we have:

Theorem 2. *Assuming p and Δ are constants such that $\left(1+\frac{1}{\Delta-1}\right)^{\Delta-1} p\Delta < 1$ (and therefore if $ep\Delta \leq 1$), there exists an integer N, which depends linearly on m, and a constant $c \in (0,1)$ (depending on p and Δ) such that if $n/\log n \geq N$ then the probability that* M-Algorithm *executes more than n calls of* Resample *is $< c^n$.*

The integer N in the above theorem is referred to as a *cut-off* point because it marks the onset of subexponential probability for the number of steps (calls of Resample) that M-Algorithm takes before it stops. Clearly, when the algorithm stops we have found an assignment such that none of the events occurs. And since this happens with probability close to 1 for large enough n, Theorem 1 indeed easily follows.

3 The acyclic chromatic index

3.1 The algorithm

Let $G = (V, E)$ be a (simple) graph with l vertices and m edges (the letter n is to indicate number of steps). The chromatic index of G is the least number of colors needed to *properly* color its edges, i.e., to color them so that no adjacent edges get the same color. If Δ is the maximum degree of G, it is known that its chromatic index is either Δ or $\Delta + 1$ (Vizing [25]).

A cycle of G of length s is a sequence $v_i, i = 0, \ldots, s-1$ of distinct vertices so that $\forall i = 0, \ldots s-1$, v_i and $v_{i+1 \pmod s}$ are connected by an edge. The *acyclic* chromatic index of G is defined as the least number of colors needed to properly color the edges of G so that no cycle is *bichromatic*, i.e., so that there is no cycle whose edges are colored with only two colors. Notice that in any properly colored graph, any cycle of odd length is necessarily at least *trichromatic*, i.e., its edges have three or more colors. It has been conjectured (J. Fiamčik [8] and Alon et al. [3]) that the acyclic chromatic index of any graph with maximum degree Δ is at most $\Delta + 2$. A number of successively tighter upper bounds to the acyclic chromatic index has been provided in the literature. Most recently, Esperet and Parreau [7] proved that the acyclic chromatic index is at most $4(\Delta - 1)$. Their proof makes use of the technique of Moser and Tardos [17] that constructively proves the Lóvasz Local Lemma. Actually, they describe a randomized algorithm and they show that it produces an acyclic edge coloring with positive probability. An approach using the entropy compression method was also used for the vertex analogue of the edge chromatic number by Gonçalves et al. [12].

In this section, we modify the Esperet and Parreau [7] technique in two important aspects: (a) instead of the Moser and Tardos [17] algorithm, we use its antecedent version by Moser [16] and (b) we analyze it by using the approach described in the previous section, which avoids the counting of witness-trees and depends only on probability estimates. Also, our technique does not refer to partial colorings. Thus we get a direct probabilistic analysis that yields the upper bound of $\lceil 3.74(\Delta-1)\rceil+1$ for the acyclic chromatic index, improving over $4(\Delta-1)$ in [7]. Generalizing to graphs with bounded girth, we also get improved numerical results some specific values of which are sampled in Figure 3. An interesting aspect of this application is that the probability of the "undesirable" events do not have a uniform upper bound, i.e., it constitutes a case of the asymmetric LLL.

Below, to facilitate notation, we call a proper edge-coloring *s-acyclic* if it contains no bichromatic cycle of length s or less (s is an even natural). We call the corresponding graph parameter the *s-acyclic chromatic index*.

We start by mentioning the following fact, proved in Esperet and Parreau [7]:

Girth	Number of colors	Previously known [7]
-	$3.731(\Delta-1)+1$	$4(\Delta-1)$
7	$3.326(\Delta-1)+1$	$3.737(\Delta-1)$
53	$2.494(\Delta-1)+1$	$3.135(\Delta-1)$
219	$2.323(\Delta-1)+1$	$3.043(\Delta-1)$

Figure 3: Our results

Lemma 7 (Esperet and Parreau [7])**.** $2(\Delta-1)+1 = 2\Delta-1$ *colors suffice to produce a 4-acyclic edge coloring of* G*.*

Proof Sketch. Successively, in any order, color the edges using, at each step, a color that does not destroy 4-acyclicity (hence, by definition, neither properness). To show that $2(\Delta-1)+1$ colors suffice, notice that for each edge e, one has to avoid the colors of all edges adjacent to e, and also for each pair of homochromatic (of the same color) edges e_1, e_2 adjacent to e at different endpoints, one has to avoid the color of the at most one edge e_3 that together with e, e_1, e_2 define a cycle of length 4. So at each step, one has to avoid $2(\Delta-1)$ colors. □

Assume now that we have $K = \lceil(2+\gamma)(\Delta-1)\rceil + 1$ colors, where γ is a nonnegative constant to be computed.

We assume below that the edges are ordered according to a fixed ordering.

Notice that the number of cycles of length $2k$ that contain a given edge e is at most $(\Delta-1)^{2k-2}$. Also, we assume that cycles are ordered according to a fixed ordering. The algorithm is given in Figure 4:

ColAlg

1: Color all edges following their ordering and choosing at each step a color uniformly at random among those that retain 4-acyclicity.
2: **while** some bichromatic cycle exists, let C be the least such and **do**
3: Recolor(C)
4: **end while**
5: Output current coloring.

Recolor *(C)*

1: Recolor the edges of C (following their ordering) choosing at each step uniformly at random a color among those that retain 4-acyclicity.
2: **while** some bicromatic cycle C' sharing an edge with C exists, let C' be the least such **do**
3: Recolor(C')
4: **end while**

Figure 4: *The coloring algorithm*

ColAlg outputs an acyclic-edge coloring *if it ever stops* (this is obvious by inspecting the algorithm). We define *root* and *recursive* calls of Recolor and *phases* in analogy to Section 2. As steps, we consider calls of Recolor.

Also we prove the following which is the analog of Lemma 1:

Lemma 8. *Consider an arbitrary call of* Recolor *(C). Let $\mathcal{E}$ be the set of edges that do not belong to a bichromatic cycle at the beginning of this call. Then, if the call terminates, the edges in $\mathcal{E}$ do not belong to any bichromatic cycle at the end of the call as well.*

Proof. Suppose some edge e'' in $\mathcal{E}$ belongs to a cycle C'' which becomes bichromatic during the call. This means that C'' shared an edge with a cycle C' such that Recolor (C') was called as a recursive call of Recolor (C) or coincides with Recolor (C). But the call Recolor (C') will not terminate until all edges in C' do not belong to a bichromatic cycle. Assuming Recolor (C) terminates, Recolor (C') must have also terminated. So C'' cannot be bichromatic at the end of Recolor (C') and therefore e'' does not belong to the bichromatic cycle C'' at the end of Recolor (C'). The same argument can be reapplied every time e'' becomes an element of a bichromatic cycle during Recolor (C). □

As an immediate corollary we get:

Corollary 2. *The cycles of root calls of* Recolor *have pairwise distinct sets of edges. Therefore the number of phases is $O(m)$ (m is the number of edges of the graph).*

In analogy to the previous section, we consider as variables the edges and as undesirable events the bichromatic cycles. Neighboring events are cycles that share an edge. We define feasible forests to be rooted forests whose nodes are labeled with cycles. We also define:

Definition 4. *A labeled rooted forest is called* feasible *if*

i. Pairwise, the labels of the roots of the trees do not share an edge,

ii. pairwise the labels of the children of every node do not share an edge and

iii. the labels of the children of every node u share an edge with the labels of u.

We define witness forests in complete analogy with the previous section. Their labels are now cycles. We define the probability $\hat{P}_n$ of ColAlg taking at least n steps. We give a validation algorithm (Figure 5) and define the probability P_n of at least one feasible forest with n nodes being validated. With respect to bounding P_n though there are some important difficulties:

(i) The edges are not colored independently.

(ii) There is no constant bound on the cardinality of neighborhoods of cycles.

(iii) The cycles are too many.

We will overcome difficulty (i) by assuming that enough colors exist as to have a margin within which colors are chosen uniformly at random. Difficulties (ii) and (iii) will be overcome by gathering together, in the computation of P_n, cycles that contain a given edge.

First, observe that:

ValColAlg (feasible forest F with root labels $C_1 \ldots, C_n$)

1: Color all edges in their ordering choosing uniformly at random a color that retains 4-acyclicity.
2: **for** $s = 1, \ldots, n$ **do**
3: **if** C_s is not bichromatic under current coloring **then**
4: **return** `failure` and exit.
5: **else**
6: Recolor the edges of C_s in their ordering choosing uniformly at
7: random a color that retains 4-acyclicity.
8: **end if**
9: **end for**
10: **return** `success`.

Figure 5: *The coloring validation algorithm*

Lemma 9. *At every instant during the execution of* ValColAlg*, the probability of assigning a specified color to a specified edge is at most* $\frac{1}{\gamma(\Delta-1)+1}$.

Proof. Notice that there exist $K = \lceil (2+\gamma)(\Delta-1) \rceil + 1$ colors. Now by Lemma 7, at most $2\Delta - 1$ colors suffice to guarantee 4-acyclicity. So, at any step of ValColAlg, we have a choice of at least $\lceil \gamma(\Delta-1) \rceil + 1$ colors to choose from uniformly at random. □

Lemma 10. *At every instant during the execution of* ValColAlg *that immediately follows a successful (i.e., no* `failure` *returned) iteration of the* **for** *loop of Step 2, the distribution of the current assignment of colors to the edges is as if* all *edges have been recolored anew. Therefore at every such instant, as well as at the beginning, just when the first assignment of colors to the edges has been completed, the probability of any cycle of length* $2k$ *being bichromatic is bounded by* $\left(\frac{1}{\gamma(\Delta-1)+1}\right)^{2k-2}$.

Proof. This is again an immediate application of the principle of Deferred Decisions, in conjunction with Lemma 9. The exponent $2k-2$ is due to the fact that the colors of one pair of consecutive edges of the cycle can be arbitrary. □

3.2 The recurrence relation

Our arguments so far are in complete analogy to Section 2. However to obtain a recurrence relation in the case of coloring, we will need to consider together the probability of cycles that share an edge being bichromatic, without resorting to union-bound as we did in Section 2 (inequality 2).

We first prove that:

Lemma 11. *The probability for two given consecutive edges* e, e' *to belong to a bichromatic cycle of length* $2k$ *(*$k \geq 3$*) is at most* $\frac{1}{\gamma(\Delta-1)} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3}$.

Proof. Let $e = \{u, v\}$ and $e' = \{v, w\}$ and assume the colors of them are "red" and "white", respectively. The probability that at least one edge $\neq e'$ stemming

from w being "red" is bounded by $1 - e^{-\frac{1}{\gamma}}$. Indeed, the probability that any specified edge $\neq e'$ stemming from w is *not* red is at least

$$1 - \frac{1}{\gamma(\Delta - 1) + 1}, \tag{9}$$

so the probability that none of them (at most $\Delta - 1$ in number) is red is at least

$$\left(1 - \frac{1}{\gamma(\Delta - 1) + 1}\right)^{\Delta - 1}. \tag{10}$$

Now we use the inequality

$$1 - x > e^{-\frac{x}{1-x}}, \forall x < 1, x \neq 0$$

(see e.g. [5, inequality (4.5.7)]), for $x = \frac{1}{\gamma(\Delta-1)+1}$, which implies, after elementary operations, that the expression in (10) is at least $e^{-\frac{1}{\gamma}}$, so the probability that at least one red edge $\neq e'$ stemming from w exists is at most $1 - e^{-\frac{1}{\gamma}}$.

We continue this argument for the $2k - 3$ successive edges of the cycle that follow e, e'. The uniquely defined last edge of the cycle has probability at most

$$\frac{1}{\gamma(\Delta - 1) + 1} < \frac{1}{\gamma(\Delta - 1)}$$

to be correctly colored (the last edge, whose both endpoints are uniquely determined, must be white). So, overall the probability that a bichromatic cycle with $2k$ edges which contains e, e' exists, is at most $\frac{1}{\gamma(\Delta-1)} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3}$, which completes the proof. □

Corollary 3. *The probability for a given edge e to belong to a bichromatic cycle of length $2k$ ($k \geq 3$) is at most $\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3}$.*

Proof. There are at most $\Delta - 1$ edges e' adjacent to e. Mutiply the result of Lemma 11 with $\Delta - 1$. □

Recall now that by the definition of a feasible forest, for each edge of G, there is at most one tree of the forest whose label contains this edge. Assume, by adding empty trees if necessary, that for every forest and for every edge, there is exactly one tree whose label contains this edge. Also, similarly, assume that from every node u of the forest labeled with a cycle C and for every edge of C there is a subtree hanging from u whose root is labeled by this edge.

Now given an edge e, let $\mathcal{F}_e$ be the family of feasible forests that are comprised of a single tree whose root label is a cycle that contains e. The we have:

Lemma 12.

$$P_n = \Pr\left[\bigcup_{F:|F|=n} V_F\right] \leq \sum_{\substack{n_1+\cdots+n_m=n \\ n_1,\ldots,n_m \geq 0}} \left(\prod_{i=1}^{m} \Pr\left[\bigcup_{T:|T|=n_i \,\&\, T \in \mathcal{F}_{e_i}} V_T\right]\right),$$

where $e_1, \ldots, e_m$ are the edges of the graph.

Proof. Indeed, fix a forest F; during the execution of ValColAlg with this specified F as input, we have by the Principle of Deferred Decisions that the assignment of colors upon the completion of a recoloring of the edges of a cycle is distributed as if all edges are colored anew (see Appendix). □

Now assume that Q_n are numbers such that for any edge e,

$$\Pr\left[\bigcup_{T:\ |T|=n\ \&\ T\in\mathcal{F}_e} V_T\right] \leq Q_n. \tag{11}$$

Then by Lemma 12 we get:

$$P_n \leq \sum_{\substack{n_1+\cdots+n_m=n\\ n_1,\ldots,n_m\geq 0}} Q_{n_1}\cdots Q_{n_m}. \tag{12}$$

We also have:

Lemma 13.

$$\Pr\left[\bigcup_{T:\ |T|=n\ \&\ T\in\mathcal{F}_e} V_T\right] \leq \sum_{k\geq 3}\left(\Pr\left[\bigcup_{\substack{C\text{is a cycle }\&\ |C|=2k\\ \&\ C\text{ contains }e}} V_C\right] \cdot \sum_{\substack{n_1+\cdots+n_{2k}=n-1\\ n_1,\ldots,n_{2k}\geq 0}} Q_{n_1}\cdots Q_{n_{2k}}\right),$$

where V_C is the event that ValColAlg *returns* `success` *on input a single node labelled with C.*

Proof. We apply union-bound to the length of the cycles that label the root of a tree T and then we argue as in the proof of Lemma 12 (see Appendix). □

By Lemma 13 and Corollary 3 we get that Q_n can be defined by the recurrence:

$$Q_n = \sum_{k\geq 3}\left(\frac{1}{\gamma}\left(1-e^{-\frac{1}{\gamma}}\right)^{2k-3} \cdot \sum_{\substack{n_1+\cdots+n_{2k}=n-1\\ n_1,\ldots,n_{2k}\geq 0}} Q_{n_1}\cdots Q_{n_{2k}}\right), \tag{13}$$

with $Q_0 = 1$.

3.3 Analytic asymptotics

We will asymptotically analyze the coefficients of the OGF $Q(z)$ of Q_n. Towards this end, multiply both sides of the first equality in (13) with z^n and sum for $n = 1,\ldots,\infty$ to get

$$Q(z) - 1 = \sum_{k\geq 3}\left[\frac{1}{\gamma}\left(1-e^{-\frac{1}{\gamma}}\right)^{2k-3} zQ(z)^{2k}\right], \tag{14}$$

with $Q(0) = 1$. Setting $W(z) = Q(z) - 1$ we get

$$W(z) = \sum_{k\geq 3} \left[\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3} z(W(z)+1)^{2k} \right], \tag{15}$$

with $W(0) = 0$. For notational convenience, set $W = W(z)$. Then from (15) we get:

$$\begin{aligned} W =& z\frac{1}{\gamma}\left(1 - e^{-\frac{1}{\gamma}}\right)^3 (W+1)^6 \\ & \cdot \frac{1}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (W+1)^2}. \end{aligned} \tag{16}$$

Set now

$$\begin{aligned} \phi(x) =& \frac{1}{\gamma}\left(1 - e^{-\frac{1}{\gamma}}\right)^3 (x+1)^6 \\ & \cdot \frac{1}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (x+1)^2}, \end{aligned} \tag{17}$$

to get from (16):

$$W = z\phi(W). \tag{18}$$

By [9, Proposition IV.5] (it is trivial to check that the hypotheses in that Theorem are satisfied for $\gamma > 0$), we obtain that , if we set

$$\rho = \frac{\phi(\tau)}{\tau}, \tag{19}$$

where τ is the (necessarily unique) solution in the interval $(0, R)$, where $R = \frac{1}{\left(1-e^{-\frac{1}{\gamma}}\right)} - 1$ is the radius of convergence of the series representing ϕ at 0, of the characteristic equation (in τ):

$$\phi(\tau) - \tau\phi'(\tau) = 0, \tag{20}$$

then $[z^n]u \bowtie \rho^n$ (i.e., $\limsup ([z^n])^{1/n} = \rho$, see [9, IV.3.2]).

Now by a simple search (through Maple, for the code see [10]) we found that for $\gamma = 1.73095$, the unique positive solution of (20) in the radius of convergence is $\tau = 0.1747094762$, and this value of τ gives to ρ in (18) $\rho = 0.9999789027 < 1$. Therefore by making use of (12) and working as in the previous section, we get:

Theorem 3. *Assuming Δ, the maximum degree of the graph G, is constant, and given the availability of at least $3.74(\Delta-1)+1$ colors, there exists an integer N, which depends linearly on m, the number of edges of G, and a constant $\rho \in (0,1)$ such that if $n/\log n \geq N$ then the probability that* ColAlg *executes more than n calls of* Recolor *is $< \rho^n$.*

Now if the graph has girth $2r - 1$ for $r \geq 4$, the previous arguments carry

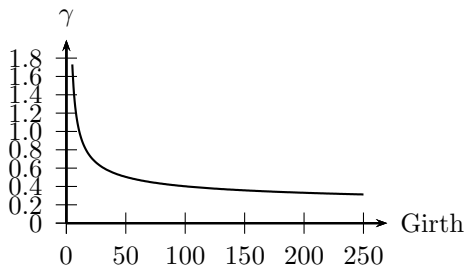

Figure 6: γ as a function of girth

over with minimal changes. Namely, equation (13) becomes:

$$Q_n = \sum_{k \geq l} \left[\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3} \sum_{\substack{n_1+\cdots+n_{2k}=n-1 \\ n_1,\ldots,n_{2k}\geq 0}} Q^*_{n_1} \cdots Q^*_{n_{2k}} \right], \quad Q_0 = 1. \tag{21}$$

Also in (14) and (15), the starting point of the summation is changed from 3 to r. Moreover, equation (16) becomes:

$$W = z \frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2r-3} (W+1)^{2r} \cdot \frac{1}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (W+1)^2} \tag{22}$$

and equation (17) becomes:

$$\phi(x) = \frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2r-3} (x+1)^{2r} \cdot \frac{1}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (x+1)^2}. \tag{23}$$

Working as before, we get numerical results depicted in Figure 6 with sample specific values explicitly given in Figure 3.

Acknowledgements

We are grateful to J. Rué for showing to us how to deal with the asymptotics of the coefficients of inverse generating functions. We are grateful to D. Mitsche on one hand and to D. Achlioptas and F. Iliopoulos, on the other, for pointing out errors in previous versions of this paper; the second and fourth authors are also indebted to the latter two for initiating them to this line of research. We are indebted to S. Messaris and Z. Terzopoulou, both undergraduate students, for their valuable comments during informal presentations of the results in this work.

References

[1] D. Achlioptas and F. Iliopoulos. Untitled notes. 2013. Unpublished private communication.

[2] Noga Alon. A parallel algorithmic version of the local lemma. *Random Structures & Algorithms*, 2(4):367–378, 1991.

[3] Noga Alon, Benny Sudakov, and Ayal Zaks. Acyclic edge colorings of graphs. *Journal of Graph Theory*, 37(3):157–167, 2001.

[4] József Beck. An algorithmic approach to the Lovász local lemma. I. *Random Structures & Algorithms*, 2(4):343–365, 1991.

[5] NIST Digital Library of Mathematical Functions. http://dlmf.nist.gov/4.5.7, Release 1.0.9 of 2014-08-29. Online companion to [18].

[6] Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Infinite and finite sets*, 10:609–627, 1975.

[7] Louis Esperet and Aline Parreau. Acyclic edge-coloring using entropy compression. *Eur. J. Comb.*, 34(6):1019–1027, 2013.

[8] J. Fiamčik. The acyclic chromatic class of a graph (in Russian). *Math. Slovaca*, 28:139–145, 1978.

[9] Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, New York, NY, USA, 1 edition, 2009.

[10] Ioannis Giotis, Lefteris Kirousis, Kostas I. Psaromiligkos, and Dimitrios M. Thilikos. Maple code. Can be found here.

[11] Ioannis Giotis, Lefteris Kirousis, Kostas I. Psaromiligkos, and Dimitrios M. Thilikos. An alternative proof for the constructive asymmetric Lovász local lemma. 2015. Presented at the 13th Twente-Cologne Workshop on Graphs & Combinatorial Optimization, May 2015, Istanbul, can be found at arXiv:1505.020568.

[12] Daniel Gonçalves, Mickaël Montassier, and Alexandre Pinlou. Entropy compression method applied to graph colorings, 2014. arXiv:1406.4380.

[13] Nicholas Harvey and Jan Vondrák. An algorithmic proof of the lopsided Lovász local lemma. arXiv:1504.02044.

[14] Donald Ervin Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*, volume 10. American Mathematical Soc., 1997.

[15] Kashyap Babu Rao Kolipaka and Mario Szegedy. Moser and Tardos meet Lováz. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 235–244, New York, NY, USA, 2011. ACM.

[16] Robin A Moser. A constructive proof of the Lovász Local Lemma. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 343–350. ACM, 2009.

[17] Robin A Moser and Gábor Tardos. A constructive proof of the general Lovász Local Lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010.

[18] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, 2010. Print companion to [5].

[19] Joel Spencer. Asymptotic lower bounds for Ramsey functions. *Discrete Mathematics*, 20(0):69 – 76, 1977.

[20] Joel Spencer. *Ten lectures on the probabilistic method*, volume 64. SIAM, 1994.

[21] Joel Spencer. Robin Moser makes Lovász Local Lemma algorithmic! 2010. Can be found here.

[22] Aravind Srinivasan. Improved algorithmic versions of the Lovász Local Lemma. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 611–620. Society for Industrial and Applied Mathematics, 2008.

[23] Mario Szegedy. The Lovász local lemma–a survey. In *Computer Science–Theory and Applications*, pages 1–11. Springer, 2013.

[24] Terence Tao. Mosers entropy compression argument. 2009. Can be found here.

[25] Vadim G Vizing. Critical graphs with a given chromatic class. *Diskret. Analiz*, 5(1):9–17, 1965.

Appendix

We describe a toy example related to Lemmata 12 and 13:

Experiment.

- Roll three fair dice. Assume that after the roll you are given (deterministically) as input a bit, 0 or 1.
- Phase 1: If the bit is 0, check the first and second dice and if any of them is an ace declare success and continue to the next phase, else exit. If the bit is 1, check the second and third dice and if any of them is an ace, continue to the next phase, else exit.
- Phase 2. Reroll the two dice examined in phase 1 and repeat the same with a new deterministic bit.

Question. What is the probability that you go through both phases successfully for at least one choice of bits?
Answer. The probability that none of the three dice is an ace is $(5/6)^3$. To succeed in both phases, for at least one choice of bits, has probability $(1 - (5/6)^3)^2 = (91/216)^2$.